arXiv:1606.03275v1 [math.ST] 10 Jun 2016

Analysis of MAP in CRP Normal-Normal model

Łukasz Rajkowski

University of Warsaw

June 13, 2016

Abstract

This article presents some asymptotic results concerning the MAP in infinite mixture models with a Chinese Restaurant process prior on partitions and a hierarchical Normal-Normal sampling scheme of observations within clusters. The limiting behaviour of the maximal posterior partition (MAP) is investigated, firstly, when the data comes from a given infinite sequence and then with the assumption that the sequence comes from iid samplings from a given distribution P. In the latter case it is proved (using the theory of Gliwienko-Cantelli set classes) that the asymptotic shape of the MAP partition is strongly associated with the problem of finding a partition of the observation space that maximises a difference between what can be considered as a 'conditional variance' of a linearly transformed P with respect to this partition and its 'entropy'. As the linear transformation depends only on the within cluster variance parameter, a consequence of the results presented is inconsistency of the MAP with iid input mechanism in the sense that even if data comes from a single-component distribution then by manipulation of the within cluster variance parameter, we obtain as the number of data increases an arbitrarily large number of clusters in the MAP. The article ends with some examples that illustrate the anomalies associated with the results obtained.

Keywords: Dirichlet process, Chinese Restaurant process, random partitions

1 Introduction

Clustering is one of the central tasks in statistical data analysis in a wide variety of applications. A Bayesian approach towards this issue is to model data as coming from a random mixture of distinguishable distributions, derive the posterior distribution of clustering related to that mixture and finally choose a clustering that maximises posterior probability. A simple implementation of this idea for n observations is the following: fix K (the suspected number of clusters), $\boldsymbol{\alpha}$ (a $\mathbb{R}^K$ vector expressing the suspected relation between sizes of the clusters), G_0 (a distribution on parameter space Θ called *the base measure*) and $\{F_\theta\}_{\theta\in\Theta}$ (a family of distributions on observation space $\mathcal{X}$ called *the component distributions*). Then sample $(p_1,\ldots,p_K)\sim \mathrm{Dir}(\boldsymbol{\alpha})$ and generate a random partition of observations by deciding for every $1\leqslant i\leqslant n$ into which cluster it should belong to; k-th cluster is chosen with probability p_k. This gives us a prior distribution on the space of partitions with at most K components. Subsequently we generate observations associated with k-th component by firstly sampling $\phi_k\sim G_0$ and then sample every observation belonging to that component independently of others from distribution F_{ϕ_k}. In case of conjugacy between G_0 and $\{F_\theta\}_{\theta\in\Theta}$, this allows us to obtain posterior distribution on the space of partitions and the problem is to compute which is the most probable – it is called the MAP partition.

There is a clear disadvantage of the solution presented above: in most cases, do not know a priori the number of clusters K. A natural way to circumvent this problem is to make this

quantity random as well, e.g. by saying that it is distributed by Poisson distribution. There is however a more subtle approach and this is by using a prior distribution on the space of all partitions, not restricted by the limit on the number of clusters. A convenient choice is the Chinese Restaurant Process: the first observation creates the first cluster and each of the subsequent observations join the existing clusters with probability proportional to their occupation or creates a new one with probability proportional to α. Since there are reasonably efficient algorithms for sampling from the posterior in such a setting (among which Algorithm 8 from [8] is the most popular), this method has gained significant popularity as a tool for assessing the clustering structure, in particular inferring about the number of clusters.

Note however that this procedure requires strong modelling assumptions: as we obtain new data the number of clusters may become arbitrarily large. Indeed it may be checked easily that the expected number of clusters generated by the Chinese Restaurant procedure grows as the logarithm of the number of observations. The behaviour of the procedure when it is misused (in the sense that there is a natural limit on the number of clusters) was investigated by J. Miller and M. Harrison in [6]. Among other things they prove in Theorem 3.4 that if the data is an independent sequence from a finite mixture distribution with t components then almost surely in the limit the posterior probability of the number of cluster being t is less than 1 and in this sense the CRP procedure is not consistent.

The inconsistency result is obtained with very general assumptions on the component distributions and the base measure. As the authors point out, in such a general setting it is very difficult to analyse the behaviour of the MAP partition. It should be noted that although the expected value of the number of clusters in the CRP process tends to infinity, it is easy to check that as far as the prior is concerned, in the limit the partition with maximal probability is the one consisting of one cluster. Therefore intuitively it might be possible that MAP partitions has reasonable properties even if it is applied to a dataset that does not satisfy some modelling assumptions.

This short article presents results that concern the asymptotic properties of the MAP partition when the base measure and component distributions are multivariate normal. In Section 2 we give formal definitions of the Chinese Restaurant Process, CRP-based model and the CRP-based Normal-Normal model in $\mathbb{R}^d$. A formula for the posterior probability on the space of partitions is also given (without the norming constant which does not affect the search of the MAP).

Section 3 includes some 'quantifying' results concerning the MAP for an infinite sequence of observations. Firstly it is proved that MAP partitions divide the set of observation into sets whose convex hulls are pairwise disjoint. This is an easy consequence of the formula for the posterior and the convexity of quadratic function. If the sequence of means of their squared norm is bounded then we obtain that the fraction of the size of the biggest cluster to the number of observation is almost surely separated from 0 (Lemma 16). The same does not hold for the size of the smallest cluster but this can be fixed if bounded support of the input distribution is assumed which is stated in Lemma 19. A simple consequence of this result is the fact that the number of components in MAP partition is almost surely bounded.

Section 4 begins with the presentation of easy consequences of results obtained in Section 3. Then an example of input discrete probability distribution is given in which almost surely we create a singleton cluster in MAP infinitely many times (Example 24) and therefore an equivalent result to Lemma 22 for the size of the minimal cluster in the MAP does not hold. In Subsection 4.1 we introduce a notion of *induced partition* which is just a decision rule of partitioning the observations according to given partition of the observation space $\mathcal{A}$. We derive a formula describing the approximate behaviour of the posterior probability of induced partition (Lemma 26) which is related to a function $\Delta(\mathcal{A})$ which is a difference between what can be considered as 'entropy' and 'conditional variance' of a linear transformation of the input distribution. In Subsection 4.2 we perform similar analysis of the MAP partition. In

order to do so we need the law of large numbers to hold almost surely for the class of all convex sets. This is ensured by Example 14 from [5] for a large class of input distributions, e.g. those that have a density with respect to Lebesgue measure. Thus we obtain that the sequences of posterior probabilities of the MAP behaves approximately as a relevant function of the family of convex hulls of the observation clusters. Moreover, it can be proved that those families are in some sense convergent to a convex P-partition of observation space, i.e. finite family of convex sets whose pairwise intersection are of P measure 0 and whose total sum is of P measure 1. It is then an easy consequence that this limiting P-partition is a maximiser of Δ function. This is summarized in Theorem 43 which states that the distance between the family of convex hulls of clusters in the MAP and the set of maximisers of Δ tends to 0.

In Section 4.3 we finish our considerations by presenting two examples of non-trivial behaviour of maximal P-partition. In case of the uniform distribution on $[-1, 1]$ it can be computed that it divides the observation space into segments of equal length, whose number depends on the in-group variance σ^2. As this parameter goes to 0, the number of clusters in the MAP may become arbitrarily large. Moreover for even number of clusters the division seems to be extremely artificial since 0 is one of the division points and one would expect a reasonable partition to distinguish between the 'center' of data and the 'outliers'. The second example is uniform distribution on unit ball: in this case for sufficiently small norm of Σ the partition into single disc does not maximise our function. Since there exist maximising partitions that are convex and in this case the set of maximisers is invariant with respect to rotations around 0 we obtain a continuous subfamily of maximising partitions. Both examples are supported by the results of numerical simulations included in the Appendix.

2 Preliminaries

Definition 1. The Chinese Restaurant Process with parameter α is a sequence of random partitions $(\mathcal{J}_n)_{n\in\mathbb{N}}$, where $\mathcal{J}_n$ is a partition of $[n] = \{1, 2, \ldots, n\}$, that satisfies

$$\mathcal{J}_{n+1} \,|\, \mathcal{J}_n = \{J_1, \ldots, J_k\} \sim \begin{cases} \{J_1, \ldots, J_i \cup \{n+1\}, \ldots, J_k\} & \text{with probability } \frac{|J_i|}{n+\alpha} \\ \{J_1, \ldots, J_k, \{n+1\}\} & \text{with probability } \frac{\alpha}{n+\alpha} \end{cases}.$$

We write $\mathcal{J}_n \sim \mathrm{CRP}(\alpha)_n$.

Remark 2. The probability weights in $\mathrm{CRP}(\alpha)_n$ are given by the formula

$$\mathbb{P}(\mathcal{J}_n = \mathcal{J}) = \frac{\alpha^{|\mathcal{J}|}}{\alpha^{(n)}} \prod_{J\in\mathcal{J}} (|J| - 1)!,$$

where $a^{(k)} = a(a+1)\ldots(a+k-1)$.

Definition 3. The CRP-based model is a Bayesian statistical model of the form

$$\begin{aligned} \mathcal{J} &\sim \mathrm{CRP}(\alpha)_n \\ \boldsymbol{\theta} = (\theta_J)_{J\in\mathcal{J}} \,|\, \mathcal{J} &\overset{\text{iid}}{\sim} G_0 \\ \mathbf{x}_J = (x_j)_{j\in J} \,|\, \mathcal{J}, \boldsymbol{\theta} &\overset{\text{iid}}{\sim} F_\theta \qquad \text{for } J \in \mathcal{J} \end{aligned}$$

where G_0 is a probability distribution on the space of parameters $(\Theta, \mathcal{F})$ and $\{F_\theta\}_{\theta\in\Theta}$ is a family of probability distributions on the space of observations $(\mathcal{X}, \mathcal{B})$, where $\mathcal{F}$ and $\mathcal{B}$ are σ-fields on Θ and $\mathcal{X}$ respectively.

Remark 4. When the F_θ's have densities f_θ with respect to some σ-finite measure ν, sampling of θ may be omitted by taking the marginal distribution of $\mathbf{x}_J$ under $\mathcal{J}$. Then the

model takes the form

$$\begin{array}{rcll} \mathcal{J} & \sim & \mathrm{CRP}(\alpha)_n & \\ \mathbf{x}_J \,|\, \mathcal{J} & \sim & F_J^{G_0} & \text{for } J \in \mathcal{J} \end{array}$$

where for $F_J^{G_0}$ is a probability distribution on $\mathcal{X}^{|J|}$ with the density

$$f_J^{G_0}(\mathbf{x}_J) = \int_\Theta \prod_{j \in J} f_\theta(x_j) \mathrm{d}G_0(\theta)$$

with respect to product measure $\nu^{|J|}$.

Definition 5. CRP-based Normal-Normal model is a CRP-based model in which $\Theta = \mathbb{R}^d$, $\mathcal{X} = \mathbb{R}^d$, $\mathcal{F}$ and $\mathcal{B}$ are σ-fields of Borel sets, $G_0 = \mathcal{N}(\mu, W^2)$ and $F_\theta = \mathcal{N}(\theta, V^2)$ for $\theta \in \Theta$, where $\mu \in \mathbb{R}^d$ and $W^2, V^2 \in \mathbb{R}^{d,d}$ are the parameters of the model.

Remark 6. Note that the usual convention of denoting the covariance matrix of multidimensional normal distribution is by using single capital letters. This is however inconsistent with the convention for one dimensional case. Since the majority of the examples presented in this article are one-dimensional, it was decided to dedicate single capital letters to the square roots of covariance matrices (in matrix sense, as in [2], Theorem 7.4).

Remark 7. In order to simplify some computations it is convenient to use the notation $[A] = A^t A$, where A is a matrix of any dimensionality. This is a bit of 'overloading'; it is the same as the notation introduced previously, $[n] = \{1, 2, \ldots, n\}$, where n is a natural number. Thorough out the text, the meaning of square brackets follows from the context.

Remark 8. The covariance matrices are symmetric, positive definite and thorough the text we will be using their properties. It is convenient to list some well known facts here; their proof is left for the appendix.

(a) the sum of symmetric positive definite matrices is symmetric positive definite.
(b) the inverse of symmetric positive definite matrix is symmetric positive definite.
(c) for each symmetric positive matrix A there exist an uniquely defined symmetric positive matrix B such that $A = [B]$. We use the notation $B = A^{1/2}$.
(d) if A, B are symmetric positive definite matrices and also $A - B$ is symmetric positive definite then $B^{-1} - A^{-1}$ is symmetric positive definite.
(e) if A, B are positive definite then $\det(A + B) \geqslant \det A$.

Remark 9. According to Remark 4, the CRP-based Normal-Normal model with $\mu = 0$ may be written in the form

$$\begin{array}{rcll} \mathcal{J} & \sim & \mathrm{CRP}(\alpha)_n & \\ \mathbf{x}_J \,|\, \mathcal{J} & \sim & F_J & \text{for } J \in \mathcal{J} \end{array}$$

where for F_J is a probability distribution on $\mathbb{R}^{d|J|}$ with the density equal to

$$f_J(\mathbf{x}_J) = \left(\frac{\det R}{\sqrt{2\pi}}\right)^{|J|} \frac{\det U}{\sqrt{|J|}\det R_{|J|}} \exp\left\{\frac{1}{2}\left(|J| \cdot \left\|R_{|J|}^{-1}[R]\overline{\mathbf{x}_J}\right\|^2 - \sum_{j \in J}\left\|Rx_j\right\|^2\right)\right\}$$

with respect to measure $\lambda^{|J|}$, where λ is a Lebesgue measure on $\mathbb{R}^d$, $R = V^{-1}$, $U = W^{-1}$, $R_m = ([R] + [U]/m)^{1/2}$ for $m \in \mathbb{N}$, $\|\cdot\|$ is standard euclidean norm in $\mathbb{R}^d$ (i.e. square root of the sum of squared coefficients) and $\overline{\mathbf{x}_J} = \frac{1}{|J|}\sum_{j \in J} x_j$. The detailed computations are left for the appendix.

Remark 10. It can be seen easily that the coefficients of the matrix R_m converge to the respective coefficients of matrix R. Therefore it follows easily that

(a) $\det R_m \to \det R$
(b) if $y_m \to y$ then $\det R_m y_m \to Ry$

Remark 11. In this paper we focus our attention on Normal-Normal models with $\mu = 0$, however it may be easily seen that this is not a real restriction. Sampling from the zero-mean Normal-Normal model and then transposing the results by vector μ is equivalent to sampling from Normal-Normal model with mean μ. Therefore all the clustering properties of the model can be investigated with the assumption that $\mu = 0$.

Remark 12. By Remark 9 together with Bayes formula, we get that the conditional probability of partition $\mathcal{J}$ in zero-mean Normal-Normal model, given the observation $\mathbf{x} = (x_j)_{j=1}^n$, is proportional to

$$\mathbb{P}(\mathcal{J} \,|\, \mathbf{x}) \propto C^{|\mathcal{J}|} \prod_{J\in\mathcal{J}} \frac{|J|!}{|J|^{3/2} \det R_{|J|}} \cdot \exp\left\{\frac{1}{2}\sum_{J\in\mathcal{J}} |J| \cdot \left\|R_{|J|}^{-1}[R]\overline{\mathbf{x}_J}\right\|^2\right\} =: Q_{\mathbf{x}}(\mathcal{J}) \tag{1}$$

where $C = \alpha \det U$. Now in order to compare the value of $\mathbb{P}(\cdot \,|\, \mathbf{x})$ for two partitions of $[n]$ it is enough to compare the values of $Q_{\mathbf{x}}(\cdot)$ for those partitions.

3 MAP partition and its basic properties

Definition 13. The maximal posterior partition (MAP) of $[n]$ with observed $\mathbf{x} = (x_i)_{i=1}^\infty$ is any partition of $[n]$ that maximises $\mathbb{P}(\cdot \,|\, \mathbf{x})$ (equivalently $Q_{\mathbf{x}}(\cdot)$). We denote it by $\hat{\mathcal{J}}(\mathbf{x})$.

Let us fix an infinite sequence $\mathbf{x} = (x_i)_{i\in\mathbb{N}}$ of vectors from $\mathbb{R}^d$. It defines the sequence $\hat{\mathcal{J}}_n = \hat{\mathcal{J}}(x_1, \ldots, x_n)$ of partitions of sets $[1], [2], \ldots$. We are interested in the behaviour of this sequence with relation to the properties of $\mathbf{x}$.

Definition 14. We say that the partition $\mathcal{J}$ of $[n]$ is *convex with respect to sequence* $(x_i)_{i=1}^n$ if sets $\operatorname{conv}\{x_j : j \in J\}$, where $J \in \mathcal{J}$, are pairwise disjoint.

Lemma 15. *For every $n \in \mathbb{N}$, $\hat{\mathcal{J}}_n$ is a convex partition with respect to $(x_i)_{i=1}^n$.*

Proof. Take any $I, J \in \hat{\mathcal{J}}_n$. Consider all partitions of $[n]$ that are obtained by replacing sets I, J with $\tilde{I}$ and $\tilde{J}$ that satisfy $|\tilde{I}| = |I|$, $|\tilde{J}| = |J|$ and $\tilde{I} \cup \tilde{J} = I \cup J$. Note that by such an operation we do not alter either the number of clusters or the size of the clusters and therefore the posterior probability of such partitions is an increasing function of

$$\begin{aligned} |I| \cdot \left\|R_{|I|}^{-1}[R]\overline{x_{\tilde{I}}}\right\|^2 + |J| \cdot \left\|R_{|J|}^{-1}[R]\overline{x_{\tilde{J}}}\right\|^2 &= \frac{1}{|I|} \cdot \left\|R_{|I|}^{-1}[R]\sum_{i\in\tilde{I}} x_i\right\|^2 + \frac{1}{|J|} \cdot \left\|R_{|J|}^{-1}[R]\sum_{j\in\tilde{J}} x_j\right\|^2 = \\ &= \frac{1}{|I|} \cdot \left\|R_{|I|}^{-1}[R]\sum_{i\in\tilde{I}} x_i\right\|^2 + \frac{1}{|J|} \cdot \left\|R_{|J|}^{-1}[R]\Big(S - \sum_{i\in\tilde{I}} x_i\Big)\right\|^2 \end{aligned} \tag{2}$$

where $S = \sum_{k\in\tilde{I}\cup\tilde{J}} x_k = \sum_{k\in I\cup J} x_k$. It may be seen quite easily that (2) defines a convex quadratic function with respect to $\sum_{i\in\tilde{I}} x_i$. We investigate its value over a finite number of possible replacements. Therefore it achieves its maximal value at the vertices of convex hull of all possible values of $\sum_{i\in\tilde{I}} x_i$. Since $\hat{\mathcal{J}}_n$ is the MAP it follows that $\sum_{i\in I} x_i$ maximises (2). Note that if there exists $i_0 \in I$ such that x_{i_0} is a convex combination of some vectors x_j for $j \in J$ then $\sum_{i\in I} x_i$ would be the same convex combination of vectors $\sum_{i\in I\setminus\{i_0\}\cup\{j\}} x_i$ for $j \in J$ and that would contradict the fact that $\sum_{i\in I} x_i$ is a vertex of the convex hull of all possible values of $\sum_{i\in\tilde{I}} x_i$. This means that the convex hulls of $\operatorname{conv}\{x_i : i \in I\}$ and $\operatorname{conv}\{x_j : j \in J\}$ are disjoint, which finishes the proof. □

If $\mathcal{J}$ is a partition, let $m(\mathcal{J}) = \min_{J\in\mathcal{J}} |J|$ and $M(\mathcal{J}) = \max_{J\in\mathcal{J}} |J|$. We now prove

Lemma 16. *If $\left(\frac{1}{n}\sum_{i=1}^n \|x_i\|^2\right)_{n=1}^\infty$ is bounded then $\liminf_{n\to\infty} M(\hat{\mathcal{J}}_n)/n > 0$.*

The proof is basing on the following technical lemma whose justification is left for the appendix.

Lemma 17. *For every $\varepsilon > 0$ there exist $K \in \mathbb{N}$ such that if $n_1, \dots, n_k \leqslant n/K$, where $n = \sum_{i=1}^{k} n_i$, then $\sqrt[n]{\prod_{i=1}^{k} n_i!/n!} < \varepsilon$.*

Proof of Lemma 16. Suppose that $\liminf M(\hat{\mathcal{J}}_n)/n = 0$. Then there exists an increasing sequence $(n_k)_{k\in\mathbb{N}}$ such that $M(\hat{\mathcal{J}}_{n_k})/n_k < 1/k$ for every $k \in \mathbb{N}$. We now prove that $\lim_{k\to\infty} \sqrt[n_k]{Q_{\mathbf{x}}(\hat{\mathcal{J}}_{n_k})/Q_{\mathbf{x}}([n_k])} = 0$, hence obtaining a contradiction. Firstly note that

$$\limsup_{k\to\infty} \sqrt[n_k]{C^{|\hat{\mathcal{J}}_{n_k}|}/C} = \limsup_{k\to\infty} C^{(|\hat{\mathcal{J}}_{n_k}|-1)/n_k} \leqslant \max\{1, C\} \tag{3}$$

By Lemma 17, it follows that, under the assumptions,

$$\lim_{k\to\infty} \sqrt[n_k]{\prod_{J\in\hat{\mathcal{J}}_{n_k}} |J|!/n_k!} = 0. \tag{4}$$

From Remark 10,

$$\limsup_{k\to\infty} \sqrt[n_k]{\frac{{n_k}^{3/2} \det R_{n_k}}{\prod_{J\in\hat{\mathcal{J}}_{n_k}} |J|^{3/2} \det R_{|J|}}} \leqslant \frac{\limsup_{k\to\infty} \sqrt[n_k]{{n_k}^{3/2} \det R_{n_k}}}{\liminf_{k\to\infty} \sqrt[n_k]{\det R^{|\hat{\mathcal{J}}_{n_k}|}}} \leqslant \frac{1}{\max\{1, \det R\}} \tag{5}$$

By the inequality between linear and quadratic means applied to every coordinate, together with Remark 8, we have

$$\sum_{J\in\hat{\mathcal{J}}_n} |J| \left\| R_{|J|}^{-1}[R]\overline{\mathbf{x}_J} \right\|^2 \leqslant \sum_{J\in\hat{\mathcal{J}}_n} \sum_{j\in J} \left\| R_{|J|}^{-1}[R]x_j \right\|^2 \leqslant \sum_{J\in\hat{\mathcal{J}}_n} \sum_{j\in J} \left\| R^{-1}[R]x_j \right\|^2 \leqslant \|R\|_2^2 \sum_{i\in[n]} \|x_i\|^2,$$

where $\|\cdot\|_2$ is a matrix norm induced by $\|\cdot\|$ (i.e. $\|A\|_2 = \sup_{\|x\|=1} \|Ax\|$). From this and assumptions of the Lemma we can easily deduce that

$$\frac{1}{n_k}\Big(\sum_{J\in\hat{\mathcal{J}}_{n_k}} |J| \left\| R_{|J|}^{-1}[R]\overline{\mathbf{x}_J} \right\|^2 - n_k \left\| R_{n_k}^{-1}[R]\overline{x_{[n_k]}} \right\|^2 \Big) \quad \text{is bounded from above.} \tag{6}$$

Gathering (3), (4), (5) and (6) together, we obtain that $\limsup \sqrt[n_k]{Q_{\mathbf{x}}(\hat{\mathcal{J}}_{n_k})/Q_{\mathbf{x}}([n_k])} = 0$. Hence, for sufficiently large n, $\mathbb{P}(\hat{\mathcal{J}}_n \,|\, \mathbf{x}) < \mathbb{P}([n] \,|\, \mathbf{x})$. This is a contradiction. □

It is now tempting to formulate a stronger result, which is $\liminf_{n\to\infty} m(\hat{\mathcal{J}}_n)/n > 0$ with the assumptions of Lemma 19. However, this is not true and this follows from Remark 18.

Remark 18. There exist a sequence $(x_n)_{n=1}^{\infty}$ such that $\left(\frac{1}{n}\sum_{i=1}^{n} \|x_i\|^2\right)_{n=1}^{\infty}$ is bounded and $\liminf_{n\to\infty} m(\hat{\mathcal{J}}_n) = 1$.

Proof. This is an easy consequence of Theorem 24 that will be introduced in the next section. □

It turns out that the following result holds

Lemma 19. *If $(x_n)_{n=1}^{\infty}$ is bounded then $\liminf_{n\to\infty} m(\hat{\mathcal{J}}_n)/n > 0$.*

Proof. Assume that $\|x_n\| \leqslant r$ for $n \in \mathbb{N}$ and suppose that $\liminf_{n\to\infty} m(\hat{\mathcal{J}}_n)/n = 0$. Let $m_n = m(\hat{\mathcal{J}}_n)$ and $M_n = M(\hat{\mathcal{J}}_n)$. By Lemma 16 we know that $\liminf_{n\to\infty} M_n/n > 0$ and hence $\liminf_{n\to\infty} m_n/M_n = 0$. We will show that for sufficiently large n, the posterior probability of $\hat{\mathcal{J}}_n$ increases if we create one cluster out of a minimal and maximal cluster. For every $n \in \mathbb{N}$ let $J_n^m, J_n^M \in \hat{\mathcal{J}}_n$ satisfy $|J_n^m| = m_n$ and $|J_n^M| = M_n$ and let $\tilde{\mathcal{J}}_n$ be a partition of $[n]$ obtained from $\hat{\mathcal{J}}_n$ by joining J_n^m with J_n^M, i.e.

$$\tilde{\mathcal{J}}_n = \hat{\mathcal{J}}_n \setminus \{J_n^m, J_n^M\} \cup \{J_n^m \cup J_n^M\}.$$

In order to clarify the notation, we omit the index n in m_n and M_n and remember that they are both functions of n. When taking a quotient $\mathbb{P}(\hat{\mathcal{J}}_n \,|\, \mathbf{x})/\mathbb{P}(\tilde{\mathcal{J}}_n \,|\, \mathbf{x})$ most factors in (1) cancel out, giving

$$\frac{\mathbb{P}(\hat{\mathcal{J}}_n \,|\, \mathbf{x})}{\mathbb{P}(\tilde{\mathcal{J}}_n \,|\, \mathbf{x})} = C\frac{m!M!}{(m+M)!}\left(\frac{m+M}{mM}\right)^{3/2}\frac{\det R_{m+M}}{\det R_m \cdot \det R_M}\cdot \exp\{D_n\}^{1/2}, \tag{7}$$

where $D_n = m\left\|R_m^{-1}[R]\overline{\mathbf{x}_{J_n^m}}\right\|^2 + M\left\|R_M^{-1}[R]\overline{\mathbf{x}_{J_n^M}}\right\|^2 - (m+M)\left\|R_{m+M}^{-1}[R]\overline{\mathbf{x}_{J_n^m\cup J_n^M}}\right\|^2$. It is now straightforward to verify that

$$\begin{aligned} m\left\|R\overline{\mathbf{x}_{J_n^m}}\right\|^2 + M\left\|R\overline{\mathbf{x}_{J_n^M}}\right\|^2 - (m+M)\left\|R\overline{\mathbf{x}_{J_n^m\cup J_n^M}}\right\|^2 = \frac{mM}{m+M}\left\|R(\overline{\mathbf{x}_{J_n^m}} - \overline{\mathbf{x}_{J_n^M}})\right\|^2 \leqslant \\ \leqslant m\left\|R(\overline{\mathbf{x}_{J_n^m}} - \overline{\mathbf{x}_{J_n^M}})\right\|^2 \leqslant m\|R\|_2^2 \cdot 4r^2. \end{aligned} \tag{8}$$

Moreover

$$\begin{aligned} (m+M)\left\|R\overline{\mathbf{x}_{J_n^m\cup J_n^M}}\right\|^2 - (m+M)\left\|R_{m+M}^{-1}[R]\overline{\mathbf{x}_{J_n^m\cup J_n^M}}\right\|^2 = \\ = (m+M)\overline{\mathbf{x}_{J_n^m\cup J_n^M}}^t R(\mathrm{I} - [R_{m+M}^{-1}R])R\overline{\mathbf{x}_{J_n^m\cup J_n^M}} = \\ = (m+M)\overline{\mathbf{x}_{J_n^m\cup J_n^M}}^t R\Big(\mathrm{I} - \big(\mathrm{I} + [UR^{-1}]/(m+M)\big)^{-1}\Big)R\overline{\mathbf{x}_{J_n^m\cup J_n^M}} = \\ = \overline{\mathbf{x}_{J_n^m\cup J_n^M}}^t R[UR^{-1}]\big(\mathrm{I} + [UR^{-1}]/(m+M)\big)^{-1}R\overline{\mathbf{x}_{J_n^m\cup J_n^M}} \end{aligned}$$

and therefore

$$\limsup_{n\to\infty}\Big((m+M)\left\|R\overline{\mathbf{x}_{J_n^m\cup J_n^M}}\right\|^2 - (m+M)\left\|R_{m+M}^{-1}[R]\overline{\mathbf{x}_{J_n^m\cup J_n^M}}\right\|^2\Big) \leqslant \|U\|_2^2 r^2. \tag{9}$$

By using Remark 8 and joining (8) and (9) we get

$$\limsup_{n\to\infty} D_n \leqslant m\|R\|_2^2 \cdot 4r^2 + \|U\|_2^2 r^2. \tag{10}$$

Using the Stirling formula, valid for every $n \in \mathbb{N}$ (cf. [4])

$$\sqrt{2\pi n}(n/e)^n e^{\frac{1}{12n+1}} < n! < \sqrt{2\pi n}(n/e)^n e^{\frac{1}{12n}}$$

we obtain

$$\frac{m!M!}{(m+M)!} \leqslant \sqrt{2\pi}\left(\frac{mM}{m+M}\right)^{1/2}\frac{m^m M^M}{(m+M)^{(m+M)}}e \leqslant \sqrt{2\pi}e\left(\frac{mM}{m+M}\right)^{1/2}\left(\frac{m}{M}\right)^m. \tag{11}$$

Now by applying (10) and (11) to (7) we obtain that for some constants C' and C''

$$\limsup_{n\to\infty}\frac{\mathbb{P}(\hat{\mathcal{J}}_n \,|\, \mathbf{x})}{\mathbb{P}(\tilde{\mathcal{J}}_n \,|\, \mathbf{x})} \leqslant \limsup_{n\to\infty} C'\frac{m+M}{mM}\left(\frac{mC''}{M}\right)^m \to 0,$$

as $m/M \to 0$. Hence, for sufficiently large n the posterior probability of $\hat{\mathcal{J}}_n$ is smaller than the posterior probability of $\tilde{\mathcal{J}}_n$. This contradicts the definition of $\hat{\mathcal{J}}_n$ and finishes the proof of the lemma. □

Corollary 20. *If $(x_n)_{n=1}^\infty$ is bounded then the sequence $(|\hat{\mathcal{J}}_n|)_{n\in\mathbb{N}}$ is also bounded.*

Proof. This follows from Lemma 19 and the obvious inequality $m(\mathcal{J}) \leqslant n/|\mathcal{J}|$. □

4 Randomisation of input

Now let P be a probability distribution on $(\mathbb{R}^d, \mathcal{B})$ and let $(X_n)_{n=1}^{\infty}$ be a sequence of independent copies of a random variable X with distribution P. Then $\hat{\mathcal{J}}_n = \hat{\mathcal{J}}(X_1, \ldots, X_n)$ becomes a random partition of $[n]$. The following results are easy consequences of Lemmas 16 and 19.

Corollary 21. *If* $\mathbb{E}\,\|X\|^4 < \infty$ *then* $\liminf_{n\to\infty} M(\hat{\mathcal{J}}_n)/n > 0$ *almost surely.*

Corollary 22. *If the support of* P *is bounded then* $\liminf_{n\to\infty} m(\hat{\mathcal{J}}_n)/n > 0$ *almost surely.*

Corollary 23. *If the support of* P *is bounded then* $|\hat{\mathcal{J}}_n|$ *is almost surely bounded.*

It turns out that there exist a probability distribution P with countable number of distant enough atoms, whose probabilities are chosen so that $\mathbb{E}\,\|X\|^4 < \infty$ and almost surely infinitely many times the last observation is a singleton in the MAP. The details are given in Theorem 24.

Theorem 24. *There exist a probability measure* P *for which* $\mathbb{E}\,\|X\|^4 < \infty$ *and almost surely* $\liminf_{n\to\infty} m(\hat{\mathcal{J}}_n) = 1$.

Proof. Take $d = 1$ and let $\alpha = V = W = 1$. By gentle modifications of (1) it is easy to show that for any partition $\mathcal{J}$ of $[n]$ if $\mathcal{J}_{\{n\}}$ is obtained by creating a singleton out of $n \in J_n \in \mathcal{J}$ and $y_1, \ldots, y_n \in \mathbb{R}^d$ then $\mathbb{P}(\mathcal{J}_{\{n\}} \,|\, y_1, \ldots, y_n)/\mathbb{P}(\mathcal{J} \,|\, y_1, \ldots, y_n)$ is equal to

$$h_{J_n}(y_1, \ldots, y_n) = \frac{1}{|J_n| - 1}\sqrt{\frac{|J_n|+1}{2|J_n|}}\exp\left\{\frac{y_n^2}{4} + \frac{(\sum y_{J_n\setminus\{n\}})^2}{2|J_n|} - \frac{(\sum y_{J_n})^2}{2(|J_n|+1)}\right\}. \tag{12}$$

Note that the content of the exponent function in the formula above is equal to

$$y_n^2\frac{|J_n|-1}{4(|J_n|+1)} - y_n\frac{\sum y_{J_n\setminus\{n\}}}{|J_n|+1} + \frac{(\sum y_{J_n\setminus\{n\}})^2}{2|J_n|(|J_n|+1)}, \tag{13}$$

which is a convex quadratic function with respect to y_n. Now, since $|J_n| \geqslant 2$, it follows easily that

$$\frac{|J_n|-1}{4(|J_n|+1)} \geqslant \frac{1}{8} \quad \text{and} \quad \left|\frac{\sum y_{J_n\setminus\{n\}}}{|J_n|+1}\right| \leqslant |\overline{y_{J_n\setminus\{n\}}}| \tag{14}$$

Now let $L = 2\cdot 9^4$ and $\tilde{x}_m = 4\cdot 9^m$. We show that if

$$n \leqslant L^{m+1}, \quad y_n \geqslant \tilde{x}_m, \quad \text{and} \quad |y_1|, \ldots, |y_{n-1}| \leqslant \tilde{x}_{m-1} \tag{$\star$}$$

then $h_{J_n}(y_1, \ldots, y_n) > 1$ (regardless of J_n) and hence in MAP partition for $[n]$ based on data $(y_i)_{i=1}^n$ singleton $\{n\}$ forms a separate cluster. Assume ($\star$). Note that if $n \leqslant L^{m+1}$ and $|y_1|, \ldots, |y_{n-1}| \leqslant \tilde{x}_{m-1}$ then by (12), (13) and (14) we obtain that

$$h_{J_n}(y_1, \ldots, y_n) \geqslant \frac{1}{L^{m+1}}\sqrt{\frac{1}{2}}\exp\left\{\frac{1}{8}y_n^2 - \tilde{x}_{m-1}y_n\right\} =: l(y_n).$$

Now as we can easily compute zeros of quadratic function, $l(y_n) \geqslant 1$ is implied by

$$y_n \geqslant 4\big(\tilde{x}_{m-1} + \sqrt{\tilde{x}_{m-1}^2 + (m+1)\ln L + (\ln 2)/2}\big).$$

It can be easily proved by induction that $\tilde{x}_m > (m+2)\ln L + (\ln 2)/2$ and therefore

$$4\big(\tilde{x}_{m-1} + \sqrt{\tilde{x}_{m-1}^2 + (m+1)\ln L + (\ln 2)/2}\big) < 9\tilde{x}_{m-1} = \tilde{x}_m$$

and as $y_n \geqslant \tilde{x}_m$ we have that $h_{J_n}(y_1, \ldots, y_n) > 1$.

Note that if $(y_n)_{n=1}^{\infty}$ is a sequence whose terms belong to $\{\tilde{x}_m \colon m \in \mathbb{N}\}$ then if for some $m \in \mathbb{N}$

$$n \leqslant L^m, \quad y_n \geqslant \tilde{x}_m, \quad \text{and} \quad y_1, \ldots, y_{n-1} < y_n. \tag{$\star'$}$$

then condition $(\star)$ holds with some $m' \geqslant m$ (the one that satisfies $\tilde{x}_{m'} = y_n$). Indeed, if $(\star')$ is satisfied and $y_n = \tilde{x}_{m'}$ then as $y_1, \ldots, y_{n-1} < y_n$ we have $y_1, \ldots, y_{n-1} \leqslant \tilde{x}_{m'-1}$, moreover $\tilde{x}_{m'} = y_n \geqslant \tilde{x}_m$ and hence $m' \geqslant m$ and $n \leqslant 2^m \leqslant 2^{m'}$ and hence $(\star)$ is satisfied.

We now prove that it is possible to choose probability weights $(p_m)_{m\in\mathbb{N}}$ so that the following probability distribution $P = \sum_{m=1}^{\infty} p_m \delta_{\tilde{x}_m}$ has a finite fourth moment and if $(X_n)_{n=1}^{\infty} \overset{\text{iid}}{\sim} P$ then $(\star')$ happens almost surely infinitely many times. Let $q = L^{-1}$ and $p_m = (1-q)q^{m-1}$. It is straightforward to check that in this case P has got finite fourth moment. Now let $s_m = \sum_{i=1}^{m} p_i = 1 - q^m$. Then $s_m^{L^m} \to e^{-1}$. Let $n_m = \sum_{i=0}^{m} L^i = \frac{L^{m+1}-1}{L-1} < L^{m+1}$. Let $A_m = \{\max_{n_{m-1}\leqslant i<n_m} X_i \geqslant \tilde{x}_m\}$. Then the probability of A_m is equal to $1 - s_{m-1}^{L^m}$ which converges to $1 - e^{-L} > \frac{1}{2}$. By the Borel-Cantelli Lemma, it follows that almost surely infinitely many of the events A_m happens. Let $(x_n)_{n=1}^{\infty}$ be a realisation of $(X_n)_{n=1}^{\infty}$ for which an infinite sequence of events $(A_{m_k})_{k=1}^{\infty}$ hold. Now let

$$\hat{n}_m = \min\{n_{m-1} \leqslant n < n_m \colon x_n = \max_{n_{m-1}\leqslant i<n_m} x_i\}.$$

Then $x_{\hat{n}_{m_k}} \geqslant \tilde{x}_{m_k}$ for $k \in \mathbb{N}$. Let $(m_{k_i})_{i=1}^{\infty}$ be a subsequence of $(m_k)_{k=1}^{\infty}$ such that $x_{n_k} < x_{n_{k_i}}$ for $k < k_i$ (such subsequence exists since $\tilde{x}_m \to \infty$). Note that $x_{\hat{n}_{m_{k_i}}} \geqslant \tilde{x}_{m_{k_i}}$, $\hat{n}_{m_{k_i}} < n_{m_{k_i}} < L^{m_{k_i}+1}$ and also

$$\text{for } l < \hat{n}_{m_{k_i}} \text{ we have } \begin{cases} x_l < x_{\hat{n}_{m_{k_i}}} & \text{if } m(l) = m_{k_i}, \\ x_l \leqslant x_{\hat{n}_{m(l)}} < x_{\hat{n}_{m_{k_i}}} & \text{if } m(l) = m_k \text{ for some } k < k_i, \\ x_l < \tilde{x}_{n_{m(l)}} < \tilde{x}_{\hat{n}_{m_{k_i}}} \leqslant x_{\hat{n}_{m_{k_i}}} & \text{otherwise,} \end{cases}$$

where $m(l) = \min\{m \in \mathbb{N} \colon n_m > l\}$. From this we may easily deduce that for every $i \in \mathbb{N}$ condition $(\star')$ is satisfied with $n = \hat{n}_{m_{k_i}}$ and $m = m_{k_i}$. This proves that almost surely the MAP partition creates a new cluster out of a new customer infinitely many times. □

4.1 Induced partitions

Definition 25. Let $\mathcal{A}$ be a partition of $\mathbb{R}^d$. For $n \in \mathbb{N}$ and $A \in \mathcal{A}$ let $J_n^A = \{i \leqslant n \colon X_i \in A\}$ and define a random partition of $[n]$ by

$$\mathcal{J}_n^{\mathcal{A}} = \left\{J_n^A \neq \emptyset \colon A \in \mathcal{A}\right\}.$$

We say that the partition of $[n]$ defined above is induced by $\mathcal{A}$.

In the following part of the text, we use the notation $p_B = P(B)$ for $B \in \mathcal{B}$ and $\mathbf{X}_{1:n} = (X_1, \ldots, X_n)$. Moreover, if $(a_n)_{n=1}^{\infty}$ and $(b_n)_{n=1}^{\infty}$ are two sequences of nonzero real numbers then

$$a_n \approx b_n \equiv \lim_{n\to\infty} a_n/b_n = 1.$$

Lemma 26. *Let $\mathcal{A}$ be a finite partition of $\mathbb{R}$ consisting of Borel sets with positive P measure. Then almost surely*

$$\sqrt[n]{Q_{\mathbf{X}_{1:n}}(\mathcal{J}^{\mathcal{A}})} \approx \tfrac{n}{e} \exp\left\{\Delta(\mathcal{A})\right\},$$

where if $\mathcal{G}$ is any finite family of measurable sets then

$$\Delta(\mathcal{G}) = \frac{1}{2}\sum_{G\in\mathcal{G}} p_G \left\|\mathbb{E}\left(RX \,|\, X \in G\right)\right\|^2 - \sum_{G\in\mathcal{G}} p_G \ln \frac{1}{p_G}.$$

Proof. We abuse the notation slightly and denote $p_{J_n^A} = |J_n^A|/n$ for $A \in \mathcal{A}$. By the law of large numbers the sequence $(X_n)_{n=1}^\infty$ almost surely satisfies $p_{J_n^A} \to p_A > 0$. By Stirling formula

$$\prod_{J\in\mathcal{J}_n^A} (np_J)! \approx \prod_{J\in\mathcal{J}_n^A} \Big(\frac{np_J}{e}\Big)^{np_J} \sqrt{2\pi np_J} = \sqrt{2\pi n}^{|\mathcal{J}_n^A|} \sqrt{\prod_{J\in\mathcal{J}_n^A} p_J} \cdot \big(\frac{n}{e} \prod_{J\in\mathcal{J}_n^A} p_J^{p_J}\big)^n \tag{15}$$

from which it follows directly that $\sqrt[n]{\prod_{J\in\mathcal{J}_n^A}(np_J)!} \approx \frac{n}{e}\prod_{J\in\mathcal{J}_n^A} p_J^{p_J}$. Note that since $J_n^{\mathcal{A}}$ has at most $|\mathcal{A}|$ elements,

$$\lim_{n\to\infty} \sqrt[n]{C^{|\mathcal{J}_n^{\mathcal{A}}|}} = 1 \quad \text{and} \quad \lim_{n\to\infty} \sqrt[n]{\prod_{J\in\mathcal{J}_n^{\mathcal{A}}} |J|^{3/2} \det R_{|J|}} = 1 \tag{16}$$

It follows from law of large numbers that $\overline{\mathbf{X}_{J_n^A}} \to \mathbb{E}\,(X \,|\, X \in A)$ for $A \in \mathcal{A}$ almost surely then (again by Remark 10)

$$\frac{1}{n}\sum_{J\in\mathcal{J}_n^{\mathcal{A}}} |J| \Big\| R_{|J|}^{-1}[R]\overline{\mathbf{X}_J}\Big\|^2 \approx \sum_{J\in\mathcal{J}_n^{\mathcal{A}}} p_J \Big\| R\overline{\mathbf{X}_J}\Big\|^2 \approx \sum_{A\in\mathcal{A}} p_A \Big\| \mathbb{E}\,(RX \,|\, X \in \mathcal{A})\Big\|^2. \tag{17}$$

Connecting (15), (16) and (17) completes the proof of the lemma. □

Remark 27. If $\mathcal{A}$ is a partition of $\mathbb{R}^d$ then $\frac{1}{2}\sum_{A\in\mathcal{A}} p_A \big\| \mathbb{E}\,(Rx_1 \,|\, x_1 \in A)\big\|$ may be considered as the expected value of the square of the conditional expected value of $Rx_1/\sqrt{2}$ conditioned on the σ-field generated by $\{x_1 \in A\colon A \in \mathcal{A}\}$. Similarly $\sum_{A\in\mathcal{A}} p_A \ln \frac{1}{p_A}$ may be interpreted as the entropy of this conditional expected value.

The following corollary is an easy consequence of Lemma 26

Corollary 28. *If $\mathcal{A}$, $\mathcal{B}$ are two partitions of $\mathbb{R}$ such that*

$$\Delta(\mathcal{A}) > \Delta(\mathcal{B})$$

then almost surely $\mathbb{P}(\mathcal{J}_n^{\mathcal{A}} \,|\, \mathbf{X}_{1:n}) > \mathbb{P}(\mathcal{J}_n^{\mathcal{B}} \,|\, \mathbf{X}_{1:n})$ for sufficiently large n.

4.2 Convergence of the MAP

Corollary 28 give us a convenient characterisation of the partitions of $\mathbb{R}^d$ that in the limit induce the best possible partitions of sets $[n]$. At this stage however we do not know yet how the best induced partitions relate to overall best partitions, namely MAPs. As we already know that clusters in MAP partition create disjoint convex sets, the analysis of the approximate behaviour of these partitions would be easier if a form of 'uniform law of large numbers' with respect to the family $\mathcal{C}$ of convex sets would be true. More precisely if we let $P_n = \frac{1}{n}\sum_{i=1}^n \delta_{X_i}$ we need the following to hold

$$\lim_{n\to\infty} \sup_{C\in\mathcal{C}} \big|P_n(C) - P(C)\big| = 0 \quad \text{almost surely.} \tag{$*$}$$

In other words we require that the class $\mathcal{C}$ is a *Glivenko-Cantelli class* with respect to P. A very convenient condition for this to hold is given in [5], Example 14:

Lemma 29. *If for each $C \in \mathcal{C}$ the boundary ∂C can be covered by countably many hyperplanes plus a set of P-measure zero, then $(*)$ holds for P.*

In particular it can be easily seen that the assumptions of 29 are satisfied if P has a density with respect to Lebesgue measure on $\mathbb{R}^d$.

Apart from approximating the probabilities of sets we also need to approximate the conditional expected value and hence we also need the following

Lemma 30. *If P has bounded support and satisfies $(*)$ then it also satisfies*

$$\lim_{n\to\infty} \sup_{C\in\mathcal{C}} \left\| \mathbb{E}_n X \mathbf{1}_{X\in C} - \mathbb{E}\, X \mathbf{1}_{X\in C} \right\| = 0 \quad \textit{almost surely.} \tag{$**$}$$

where $E_n f(X) = \int_{\mathcal{X}} f(X)\mathrm{d}P_n = \frac{1}{n}\sum_{i=1}^n f(X_i)$.

Proof. Let $x^{(i)}$ $(r \leqslant d)$ be i-th coordinate of vector x. We now prove that

$$\lim_{n\to\infty} \sup_{C\in\mathcal{C}} \left| \mathbb{E}_n X^{(1)} \mathbf{1}_{X\in C} - \mathbb{E}\, X^{(1)} \mathbf{1}_{X\in C} \right| = 0. \tag{18}$$

By proving the same thing for the remaining coordinates, $(**)$ follows directly. Assume that supp $P \subseteq [-r,r]^d$. Now fix any $C \in \mathcal{C}$ and for $m \in \mathbb{N}$ and $-m \leqslant k \leqslant m-1$ let $C_k^m = C \cap [rk/m, r(k+1)/m) \times \mathbb{R}^{d-1}$. Then

$$\left| \mathbb{E}\, X^{(1)} \mathbf{1}_{X\in C} - \sum_{k=-m}^{m-1} r\frac{k}{m} P(C_k^m) \right| \leqslant \frac{r}{m} P(C) \leqslant \frac{r}{m}. \tag{19}$$

It follows from the same reasoning

$$\left| \mathbb{E}_n X^{(1)} \mathbf{1}_{X\in C} - \sum_{k=-m}^{m-1} r\frac{k}{m} P_n(C_k^m) \right| \leqslant \frac{r}{m} \quad \text{for every } n \in \mathbb{N}. \tag{20}$$

Now choose $\varepsilon > 0$ and $m > r/\varepsilon$. Note that C_k^m are convex sets (as intersections of two convex sets) and hence by $(*)$ we may choose N so that for $n > N$ and any convex C' we have that $|P_n(C) - P(C)| < \varepsilon/(m+1)$ and hence

$$\left| \sum_{k=-m}^{m-1} r\frac{k}{m} P(C_k^m) - \sum_{k=-m}^{m-1} r\frac{k}{m} P_n(C_k^m) \right| \leqslant \sum_{k=-m}^{m-1} r\frac{|k|}{m} |P(C_k^m) - P_n(C_k^m)| < r\varepsilon. \tag{21}$$

By combining (19), (20) and (21) we obtain that $|\mathbb{E}_n X^{(1)} \mathbf{1}_{X\in C} - \mathbb{E}\, X^{(1)}| \mathbf{1}_{X\in C} < (2+r)\varepsilon$ for $n > N$ and since the choice of N does not depend on C, (18) follows, finishing the proof of the Lemma. □

Lemma 31. *Assume that P has a bounded support and satisfies $(*)$ and let $\hat{\mathcal{A}}_n = \left\{ conv\{\mathbf{X}_j : j \in J\} : J \in \hat{\mathcal{J}} \right\}$. Then almost surely*

$$\sqrt[n]{Q_{\mathbf{X}_{1:n}}(\hat{\mathcal{J}}_n)} \approx \frac{n}{e} \exp\{\Delta(\hat{\mathcal{A}}_n)\}.$$

Proof. From Lemma 21 we know $\min\{np_J : J \in \hat{\mathcal{J}}_n\} \to \infty$. By applying Stirling formula to each factor $(np_J)!$ and taking into account that the number of factors is bounded, we obtain that

$$\frac{\prod_{J\in\hat{\mathcal{J}}_n} (np_J)!}{n!} \approx \frac{\prod_{J\in\hat{\mathcal{J}}_n} \left(\frac{np_J}{e}\right)^{np_J} \sqrt{2\pi np_J}}{\left(\frac{n}{e}\right)^n \sqrt{2\pi n}} = \sqrt{2\pi n}^{|\hat{\mathcal{J}}_n|-1} \sqrt{\prod_{J\in\hat{\mathcal{J}}_n} p_J} \cdot \left(\prod_{J\in\hat{\mathcal{J}}_n} p_J^{p_J} \right)^n.$$

Obviously the elements of $\hat{\mathcal{A}}_n$ are convex and hence by Lemma 29 the frequencies p_J for $J \in \mathcal{J}_n$ uniformly approximate respective probabilities of sets in $\hat{\mathcal{A}}_n$. Hence, as $(|\hat{\mathcal{J}}_n|)_{n=1}^\infty$ is almost surely bounded, it follows that

$$\sqrt[n]{\frac{\prod_{J\in\hat{\mathcal{J}}_n} (np_J)!}{n!}} \approx \prod_{J\in\hat{\mathcal{J}}_n} p_J^{p_J} \approx \prod_{A\in\mathcal{A}_n} p_A^{p_A}.$$

By applying a similar argument to the remaining part of formula (1), the result follows by Lemma 30. □

Now we investigate the convergence of the sequence $\hat{\mathcal{A}}_n$ defined in Lemma 31. In order to do so we need a topology on some subspaces of $2^{\mathbb{R}}$. We begin by recalling two standard metrics used in this context.

Definition 32. Let $\mathcal{D}$ be a class of closed subsets of $\mathbb{R}^d$. Then the function $d_H \colon \mathcal{C}^2 \longrightarrow \mathbb{R}$ defined by

$$d_H(A,B) = \inf\{\varepsilon > 0 \colon A \subseteq (B)_\varepsilon, B \subseteq (A)_\varepsilon\},$$

where $(X)_\varepsilon = \{x \in \mathbb{R}^d \colon \operatorname{dist}(x,X) < \varepsilon\}$, is a metric on $\mathcal{D}$. It is called *Hausdorff distance.*

Definition 33. Let $\mathcal{M}$ be a σ-field and μ be a measure. Then the function $d_\mu \colon \mathcal{M}^2 \longrightarrow \mathbb{R}$ defined by

$$d_\mu(A,B) = \mu(A \div B),$$

where $A \div B = (A \setminus B) \cup (B \setminus A)$, is a pseudometric on $\mathcal{M}$ i.e. it is symmetric, nonnegative and satisfies triangle inequality. It is called *the symmetric difference metric.*

Let $\mathcal{K} = \mathcal{C} \cap \mathcal{D}$ (a family of convex and closed sets). The following theorem is Theorem 3.2.14 in [7].

Theorem 34. *The space $(\mathcal{K}, d_H)$, where $\mathcal{K} = \mathcal{D} \cap \mathcal{C}$, is finitely compact (i.e. every bounded sequence has a convergent subsequence).*

It turns out that Hausdorff metric restricted to a bounded subspace of $\mathcal{K}$ is weaker than symmetric difference metric. This is formally stated in Theorem 35.

Theorem 35. *In the space $\mathcal{K}_r = \mathcal{K} \cap B(0,r)$ convergence in d_H implies convergence in d_P.*

Proof. We prove that if $K_n \in \mathcal{K}_r$ and $d_H(K_n, K_0) \to 0$ then $d_P(K_n, K_0) \to 0$. Since P is absolutely continuous with respect to Lebesgue measure λ it suffices to show that $d_\lambda(K_n, K) \to 0$. Now let V_k, $k = 0, \ldots, d$, be Steiner functionals (the definition may be found in [7], Chapter 7). By Theorem 7.2.12 in [7] we get that for $K \in \mathcal{K}_r$ and $\alpha > 0$

$$\lambda(\operatorname{col}_\alpha K) = \alpha \sum_{k=0}^{n-1} \alpha^{n-k-1} \kappa_{n-k} V_k(K)$$

(where $\operatorname{col}_\alpha K = (K)_\alpha \setminus K$ and κ_i is a volume of i-dimensional unitary ball), which by monotonicity of Steiner functionals (cf. [7], Corollary 7.2.11) give us

$$\lambda(\operatorname{col}_\alpha K) \leqslant \alpha M \quad \text{where } M = \sum_{k=0}^{n-1} \alpha^{n-k-1} \kappa_{n-k} V_k(B(0,r)).$$

Now fix $\varepsilon > 0$; as $d_H(K_n, K_0) \to 0$ there exists $N \in \mathbb{N}$ such that for $n \geqslant N$ we have

$$K_n \subseteq (K_0)_\varepsilon \quad \text{and} \quad K_0 \subseteq (K_n)_\varepsilon$$

which means that

$$K_n \setminus K_0 \subseteq \operatorname{col}_\varepsilon K_0 \quad \text{and} \quad K_0 \setminus K_n \subseteq \operatorname{col}_\varepsilon K_n$$

and hence $\lambda(K_n \div K_0) \leqslant 2\varepsilon M$ for $n \geqslant N$, which finishes the proof. □

Corollary 36. *Space $(\mathcal{K}_r, d_P)$ is compact.*

Definition 37. Let d be a (pseudo)metric on the family of sets $\mathcal{F}$. Let $\mathcal{A} = \{A^{(1)}, \ldots, A^{(k)}\}$ and $\mathcal{B} = \{B^{(1)}, \ldots, B^{l}\}$ be finite subsets of $\mathcal{F}$ and without loosing generality assume that $k < l$. Let $A^{(k+1)}, \ldots, A^{(l)} = \emptyset$. We define

$$\bar{d}(\mathcal{A}, \mathcal{B}) = \min_{\sigma \in \Sigma_l} \max_{i \leqslant l} d(A^{(i)}, B^{(\sigma(i))}),$$

where Σ_n is the set of all permutations of $[n]$.

Definition 38. If $\mathcal{F}$ is a family of sets let $F_K(\mathcal{F})$ be the space of finite subfamilies of $\mathcal{F}$ that have at most K elements.

Remark 39. If $(\mathcal{F}, d)$ is a (pseudo)metric space then $(F_K(\mathcal{F}), \bar{d})$ is also a (pseudo)metric space. The proof is straighforward and therefore it is omitted.

Theorem 40. *If $(\mathcal{F}, d)$ is finitely compact then $(F_K(\mathcal{F}), \bar{d})$ is also finitely compact.*

Proof. Let $(\mathcal{A}_n)_{n=1}^{\infty}$ be a sequence in $F_K(\mathcal{F})$ and let $\mathcal{A}_n = \{A_n^{(1)}, A_n^{(2)}, \ldots, A_n^{(k_n)}\}$. As the sequence $(k_n)_{n=1}^{\infty}$ is bounded by K we may choose a subsequence $\mathcal{A}_{n_k}$ and $\tilde{K} \in \mathbb{N}$ such that $|\mathcal{A}_{n_k}| = \tilde{K}$ for every $k \in \mathbb{N}$. Consider the sequence $(A_{n_k}^{(1)})_{k=1}^{\infty}$. This sequence is bounded (as $(\mathcal{A}_n)_{n=1}^{\infty}$ is bounded). Therefore it has a subsequence $(A_{n_{k_l}}^{(1)})_{l=1}^{\infty}$ converging in d to $A^{(1)} \in \mathcal{F}$. Now we consider $(A_{n_{k_l}}^{(2)})_{l=1}^{\infty}$ and again we choose a subsequence $(A_{n_{k_{l_m}}})_{m=1}^{\infty}$ converging in d to $A^{(2)} \in \mathcal{F}$. By iterating this procedure $\tilde{K}$ times we obtain a family $\hat{A} = \{A^{(1)}, \ldots, A^{(\tilde{K})}\}$ of 'limiting' sets. It is easy to verify that final subsequence of $(\mathcal{A}_n)_{n=1}^{\infty}$ converges in $\bar{d}$ to $\hat{A}$, which finishes the proof. □

Now assume that P has bounded support and satisfies $(*)$. Then by Corollary 36 and Theorem 40 we obtain that $(\hat{\mathcal{A}}_n)_{n=1}^{\infty}$ has a nonempty set of limits of its subsequences in $\bar{d}_P$. Let us denote this (random) set by $\boldsymbol{E}$.

Definition 41. We say that a family $\mathcal{A}$ of P-measurable sets is a *P-partition of* $\mathbb{R}^d$ if $P(\bigcup \mathcal{A}) = 1$ and $P(A_1 \cap A_2) = 0$ for every $A_1, A_2 \in \mathcal{A}$.

Theorem 42. *Every partition in* $\mathbf{E}$ *is a P-partition that maximises Δ.*

Proof. Take any $\mathcal{E} = (E^{(1)}, \ldots, E^{(\tilde{K})}) \in \mathbf{E}$ and assume that it is a limit of $\hat{\mathcal{A}}_{n_k}$ in the sense described before. The fact $P(E^{(i)} \cap E^{(j)}) = 0$ for $1 \leqslant i < j \leqslant k$ is an easy consequence of continuity of the union and intersection with respect to d_P (cf. [3], formulas (13.3)). In order to prove that $\mathcal{E}$ is P-partition it is left to show that $P(\bigcup \mathcal{E}) = 1$. Suppose this is not the case; then $E_0 = \mathbb{R}^d \setminus \bigcup \mathcal{E}$ is an open set with a positive probability. Hence it includes a ball B' of positive probability and as almost surely $p_{J_n^{B'}} \to p_{B'} > 0$ there exist $x_{n'} \in B'$. This is however not possible as $x_{n'} \in \bigcup \mathcal{E}_{n'} \subseteq \bigcup \mathcal{E}$ and hence $\mathcal{E}$ is a P-partition.

Note that by Lemma (31) and an obvious continuity of Δ with respect to $\bar{d}_P$ metric we obtain that

$$\sqrt[n]{Q_{\mathbf{X}_{1:n}}(\hat{\mathcal{J}}_n)} \approx \exp(\Delta(\hat{\mathcal{A}}_n)) \approx \exp(\Delta(\mathcal{E})). \tag{22}$$

Now take any finite P-partition $\mathcal{A}$; note that it may be assumed that every X_n belongs to exactly one of the sets in $\mathcal{A}$, $p_{\mathcal{J}_n^A} \to p_A$ and $\overline{\mathbf{X}_{\mathcal{J}_n^A}} \to \mathbb{E}\,(X \mid X \in A)$ for $A \in \mathcal{A}$ (it just requires adding a countable number of conditions on the infinite iid sequence with distribution P, each of which is satisfied almost surely). By definition of $\hat{\mathcal{J}}_n$ and Lemma 26 we get

$$\sqrt[n]{Q_{\mathbf{X}_{1:n}}(\hat{\mathcal{J}}_n)} \geqslant \sqrt[n]{Q_{\mathbf{X}_{1:n}}(\mathcal{J}_n^{\mathcal{A}})} \approx \exp(\Delta(\mathcal{A})). \tag{23}$$

Equations (22) and (23) together give us $\Delta(\mathcal{E}) \geqslant \Delta(\mathcal{A})$ which proves that $\mathcal{E}$ is a finite partition that maximises Δ. □

Let $\boldsymbol{M}_\Delta$ be the set of P-partitions that maximise Δ. It is an intersection of the set of maximisers of $\boldsymbol{M}_\Delta$ in $(F_K(\mathcal{K}_r), \bar{d}_P)$ and the subspace of P-partitions, both of them being closed subspaces of $(F_K(\mathcal{K}_r), \bar{d}_P)$. We already know that $\boldsymbol{E} \subseteq \boldsymbol{M}_\Delta$. By using an easy topological Lemma 49 presented in the Appendix, we may now deduce the following Theorem

Theorem 43. *Almost surely* $\operatorname{dist}(\hat{\mathcal{A}}_n, \boldsymbol{M}_\Delta) \to 0$ *in* $(F_K(\mathcal{K}_r), \bar{d}_P)$.

Remark 44. Theorem 43 implies that limits of MAP partitions of observations are P-partitions that maximise the function Δ. It is worth noting that this function does not depend neither on the concentration parameter α nor covariance matrix $[T]$ and only on R which is an inverse of Σ. Therefore in the limit α and T parameters does not affect the behaviour of MAP partition. Moreover it can be easily seen that as the norm of within group covariance matrix tends to 0, the variance of the conditional expected value gains larger importance in maximising Δ function and this variance gets larger as the number of clusters increase. Therefore by manipulating the within group covariance parameter it is possible to obtain arbitrarily large (but fixed) number of clusters in the MAP for sufficiently many observations.

4.3 Examples

Example 45. We now find the convex partition that maximises Δ if P is a uniform distribution on $[-1, 1]$. Since it is convex it is defined by the lengths of consecutive subsegments of $[-1, 1]$; let those be $2p_1, \ldots, 2p_n$. Let $s_k = \sum_{i=1}^{k} p_i$ for $k \geqslant 1$ and $A_k = [s_{k-1}, s_k]$, where $s_0 = 0$. Then it follows that the optimal partition maximises

$$F(p_1, \ldots, p_n) = \rho \sum_{i=1}^{n} p_i (2s_{i-1} - 1 + p_i)^2 + \sum_{i=1}^{n} p_i \ln p_i,$$

where $\rho = r^2/2$, with the constraint $\sum_{i=1}^{n} p_i = 1$. This problem can be solved using Lagrange multipliers. We are looking for the local maximum of a function

$$F_\lambda(p_1, \ldots, p_n) = F(p_1, \ldots, p_n) - \lambda \sum_{i=1}^{n} p_i.$$

We now compute its partial derivatives

$$\begin{aligned}
\tfrac{\partial}{\partial p_k} \textstyle\sum_{i=1}^{n} p_i (2s_{i-1} - 1 + p_i)^2 \quad &= \tfrac{\partial}{\partial p_k} p_k (2s_{k-1} - 1 + p_k)^2 + \tfrac{\partial}{\partial p_k} \textstyle\sum_{i=k+1}^{n} p_i (2s_{i-1} + p_i)^2 = \\
&= (2s_{k-1} - 1 + p_k)^2 + p_k \cdot 2(2s_{k-1} - 1 + p_k) + \\
&\quad + \textstyle\sum_{i=k+1}^{n} p_i \cdot 4(2s_{i-1} - 1 + p_i) = \\
&= (2s_{k-1} - 1 + p_k)^2 - 2p_k \cdot (2s_{k-1} - 1 + p_k) + \\
&\quad + 4 \textstyle\sum_{i=k}^{n} p_i (2s_{i-1} - 1 + p_i) = \\
&= (2s_{k-1} - 1)^2 - p_k^2 + 4 \textstyle\sum_{i=k}^{n} p_i (2s_{i-1} - 1 + p_i) = \\
&= (2s_n - 1)^2 - p_k^2 = 1 - p_k^2
\end{aligned}$$

from which

$$\frac{\partial}{\partial p_k} F_\lambda(p_1, \ldots, p_k) = \ln p_k - \rho p_k^2 + 1 + \rho - \lambda.$$

If all partial derivatives are zero then $\ln p_k - \rho p_k^2 = C$ for all $1 \leqslant k \leqslant n$ and some $C \in \mathbb{R}$. Therefore we may restrict the search of maximum of F on the set of probability weights to the subset where $\ln p_k - \rho p_k^2 = C$. On that set function $\tilde{F}$ is equal to

$$\tilde{F}(p_1, \ldots, p_n) = \rho \sum_{i=1}^{n} p_i (2s_{i-1} - 1 + p_i)^2 + \sum_{i=1}^{n} p_i (C + p_i^2)$$

and the derivative of this function is equal to

$$\frac{\partial}{\partial p_k} \tilde{F}(p_1, \ldots, p_n) = \rho(1 - p_k^2) + 3Cp_k^2 + C.$$

If we apply Lagrange multipliers to the function $\tilde{F}$ we get that $p_1 = p_2 = \ldots = p_n = 1/n$. Here we also have the maximum of F on the set of probability weights.

For $p_1 = p_2 = \ldots = p_n = 1/n$ we have

$$\begin{aligned}\mathbb{E}\left(\mathbb{E}\left(X \mid \mathcal{A}\right)\right)^2 &= \tfrac{1}{n}\sum_{i=1}^{n}(2\tfrac{i-1}{n} - 1 + \tfrac{1}{n})^2 = \\ &= \tfrac{1}{n^3}\sum_{i=1}^{n}\left(4i^2 - 4i(n+1) + (n+1)^2\right) = \\ &= \tfrac{4}{n^3}\tfrac{n(n+1)(2n+1)}{6} - \tfrac{4(n+1)}{n^3}\tfrac{n(n+1)}{2} + \tfrac{1}{n^3}n(n+1)^2 = \\ &= \tfrac{n+1}{n^2}\left(\tfrac{2}{3}(2n+1) - 2(n+1) + (n+1)\right) = \tfrac{(n+1)(n-1)}{3n^2} = \tfrac{1}{3}\left(1 - \tfrac{1}{n^2}\right)\end{aligned}$$

and hence

$$f(n) := F(1/n, \ldots, 1/n) = \frac{\rho}{3}\left(1 - \frac{1}{n^2}\right) - \ln n = \frac{\rho}{3} - \frac{\rho}{3n^2} - \ln n.$$

The derivative of the function $\frac{\rho}{3x^2} + \ln x$ is $-\frac{2\rho}{3x^3} + \frac{1}{x}$ so it is increasing for $x < \sqrt{\frac{2\rho}{3}} = 1/\sqrt{3\sigma^2}$ and decreasing for $x > 1/\sqrt{3\sigma^2}$ and therefore $f(n)$ achieves its maximum for $\lfloor 1/\sqrt{3\sigma^2} \rfloor$ or $\lceil 1/\sqrt{3\sigma^2} \rceil$.

Remark 46. It should be underlined that Example 45 presents a very undesirable property of MAP partition. Firstly the number of clusters in the MAP becomes arbitrarily large as σ^2 goes to 0 and it is difficult to propose any reasonable lower bound on this parameter basing on data only (as far as the upper bound is concerned, the variance of data would do). Secondly the partitioning mechanism itself is not satisfactory as for the even number of clusters it does not divide data into 'center' and the outliers but it puts a separator even in the mean of the dataset which is 0. An illustration of such peculiar division is shown on Figure 2.

Example 47. Let $B_1 = B(0,1)$ be a disk in $\mathbb{R}^2$ of radius 1 and P be uniform distribution on B_1. It can be easily seen that $\Delta(\{B_1\}) = 0$. Let $R = r\mathrm{I}$. Then if B_1^+ (B_1^-) is a subset of B_1 with non negative (negative) first coordinate then it might be easily computed that $\Delta(\{B_r^+, B_r^-\}) = 2r^2/9 - \ln 2$. Therefore for sufficiently large r partitioning into halves is better than one cluster. Hence the (existing) optimal convex partition $\mathcal{E}$ is not a single cluster. Moreover it is easy to see that a single cluster is the only convex partition that is invariant with respect to all rotations and since P is invariant with respect to all rotations we can see that rotation of $\mathcal{E}$ gives a family of optimal partitions. An approximation of the MAP partition on a disc is presented on Figure 3.

5 Discussion

As shown by Example 45, the application of the MAP procedure in CRP Normal-Normal models should be restricted to cases where the modelling assumptions are consistent with the character of the Chinese Restaurant Process (e.g. there is a sort of 'rich gets richer' behaviour). Applying it blindly to an arbitrary classification problem may yield significant overestimation of the number of clusters.

It should be also noticed that this apparent inconsistency of MAP in CRP model is in a weak sense, namely the model is misspecified not only in terms of the assumption on data generation mechanism but also as far as the component distributions are concerned (we allowed manipulation of the within group covariance parameter). It would be would be very informative to provide similar analysis in case that within group covariance matrix is properly defined. A more difficult thing would be to examine the case where P (the input distribution) is also normally distributed. A distant aim of this analysis would be to perform the analysis of the MAP in the setting of similar level of generality to the work of Miller and Harrison with the same assumption that input distribution P belongs to the family of component distributions.

However, the machinery presented can be used for a different task. The P-partitions that maximise the Δ function seem to be interesting objects in their own right. Note that for

dimension greater than 1 it seems to be extremely difficult to derive the shape of the maximisers simply by analytical means. Theorems 40 and 43 give us a convenient tool to examine those maximisers as they may be approximated by performing MCMC sampling. Further examination of the shape of the maximisers and the effectiveness of the MCMC approximation is left for future research.

6 Appendix

6.1 Proofs left for the Appendix

Remark 4 According to Remark 4, zero-mean CRP-based Normal-Normal model may be written in the form

$$\begin{array}{rcll} \mathcal{J} & \sim & \mathrm{CRP}(\alpha)_n & \\ \mathbf{x}_J \,|\, \mathcal{J} & \sim & F_J & \text{for } J \in \mathcal{J} \end{array}$$

where for F_J is a probability distribution on $\mathbb{R}^{d|J|}$ with the density equal to

$$f_J(\mathbf{x}_J) = \Big(\frac{\det R}{\sqrt{2\pi}}\Big)^{|J|} \frac{\det U}{\sqrt{|J|}\det R_J} \exp\Big\{\frac{1}{2}\Big(|J|\cdot\big\|R_J^{-1}[R]\overline{\mathbf{x}_J}\big\|^2 - \sum_{j\in J}\big\|Rx_j\big\|^2\Big)\Big\}$$

with respect to measure $\lambda^{|J|}$, where λ is a Lebesgue measure on $\mathbb{R}^d$, $R = \Sigma^{-1}$, $U = \mathrm{T}^{-1}$ and R_J is symmetric positive definite matrix such that $[R_J] = [R] + [U]/|J|$.

Proof. Since the densities of θ and $\mathbf{x}_J \,|\, \theta$ are given by

$$\theta \sim \frac{\det U}{\sqrt{2\pi}^d}\exp\Big\{-\frac{1}{2}[U\theta]\Big\} \quad \text{and} \quad \mathbf{x}_J \,|\, \theta \sim \Big(\frac{\det R}{\sqrt{2\pi}^d}\Big)^{|J|}\exp\left\{-\frac{1}{2}\sum_{j\in J}[R(x_j-\theta)]\right\}$$

we obtain that

$$f_J(\mathbf{x}_J) = \frac{\det U}{\sqrt{2\pi}^d}\left(\frac{\det R}{\sqrt{2\pi}^d}\right)^{|J|}\int_\Theta \exp\left\{-\frac{1}{2}[U\theta] - \frac{1}{2}\sum_{j\in J}[R(x_j-\theta)]\right\}\mathrm{d}\theta. \tag{24}$$

Let H_J be a positive definite symmetric matrix such that $[H_J] = |J|\cdot[R] + [U]$ and

$$\begin{aligned}[U\theta] + \sum_{j\in J}[R(x_j-\theta)] &= [U\theta] + \Big(\sum_{j\in J}[Rx_j] - 2\theta^t[R]\sum_{j\in J}x_j + |J|[R\theta]\Big) = \\ &= [H_J\theta] - 2\theta^t[R]\sum_{j\in J}x_j + \sum_{j\in J}[Rx_j] = \\ &= \Big[H_J\big(\theta - [H_J]^{-1}[R]\sum_{j\in J}x_j\big)\Big] - \Big[H_J^{-1}[R]\sum_{j\in J}x_j\Big] + \sum_{j\in J}[Rx_j]\end{aligned}$$

and hence

$$\int_\Theta \exp\left\{-\frac{1}{2}[U\theta] - \frac{1}{2}\sum_{j\in J}[R(x_j-\theta)]\right\}\mathrm{d}\theta = \frac{\sqrt{2\pi}^d}{\det H_J}\exp\left\{\frac{1}{2}\Big(\Big[H_J^{-1}[R]\sum_{j\in J}x_j\Big] - \sum_{j\in J}[Rx_j]\Big)\right\}. \tag{25}$$

By joining equalities (24) with (25) and substituting $H_J = \sqrt{|J|}\cdot R_J$ we obtain desired result. □

Remark 8 Properties of symmetric, positive definite matrices.

(a) the sum of symmetric positive definite matrices is symmetric positive definite.

(b) the inverse of symmetric positive definite matrix is symmetric positive definite.

(c) for each symmetric positive matrix A there exist an uniquely defined symmetric positive matrix B such that $A = [B]$. We use the notation $B = A^{1/2}$.

(d) if A, B are symmetric positive definite matrices and also $A - B$ is symmetric positive definite then $B^{-1} - A^{-1}$ is symmetric positive definite.

(e) if A, B are positive definite then $\det(A + B) \geqslant \det A$.

Proof. Let $A, B \in \mathbb{R}^{d,d}$.

(a) If A, B are symmetric then $A + B$ is also symmetric. If A, B are positive definite then for every $x \in \mathbb{R}^d \setminus \{\vec{0}\}$ we have $x^t(A + B)x = x^tAx + x^tBx > 0$ and hence $A + B$ is also positive definite.

(b) If A is symmetric then A^{-1} is also symmetric. If A is positive definite then by Theorem 7.1 from [2] it may be expressed as $U^*\mathrm{diag}(\lambda_1, \ldots, \lambda_d)U$ where U is unitary matrix and U^* its conjugate transpose and $\lambda_1, \ldots, \lambda_d > 0$. Therefore $A^{-1} = U^*\mathrm{diag}(\lambda_1^{-1}, \ldots, \lambda_d^{-1})U$ and again by using Theorem 7.1 we obtain that A^{-1} is positive definite.

(c) Since if A is a symmetric matrix then $[A] = A^2$ this point is an easy consequence of Theorem 7.4 in [2].

(d) Let P be a symmetric matrix that satisfy $P^2 = B$. Positive definiteness of $A - B$ is equivalent to $x'Ax > x'Bx$ for all $x \in \mathbb{R}^d$. By substituting $y = Px$ this is equivalent to $y'P^{-1}AP^{-1}y > y'y$ for all $y \in \mathbb{R}^d$. Note that $P^{-1}AP^{-1}$ is positive definite (as a product of positive definite matrices) and hence it can be expressed as $U^*\Lambda U$. But then the latest condition can be expressed as $z'z > z'U^*Lambda^{-1}Uz$ for all $z \in \mathbb{R}^d$ which in the same way is equivalent to the potive definiteness of $B^{-1} - A^{-1}$.

(e) This is obviously equivalent to $\det(\mathrm{I} + BA^{-1}) \geqslant \det(\mathrm{I}) = 1$. As BA^{-1} is positive definite then for every eigenvalue λ of $\mathrm{I} + BA^{-1}$ we have $\lambda = v'(\mathrm{I} + BA^{-1})v > 1$, where v is its eigenvector of norm 1. Therefore the determinant of $\mathrm{I} + BA^{-1}$ is also greater than 1.

□

Lemma 48. *If $n_1, \ldots, n_k \in \mathbb{N}$, $n_1 \leqslant n_2 \leqslant \ldots \leqslant n_k$ and $n = \sum_{i=1}^k n_i = an_k + r$, where $a \in \mathbb{N}$, $r < n_k$. Then $\prod_{i=1}^k n_i! \leqslant (n_k!)^a n_k(n_k - 1)\ldots(n_k - r + 1)$.*

Proof. We prove by induction on n_k that the sequence

$$[n_k]_{a,r} = (\underbrace{1, \ldots, n_k, 1, \ldots, n_k, \ldots, 1, \ldots, n_k}_{a}, n_k - r + 1, n_k - r + 2, \ldots, n_k)$$

may be ordered so that it is term-wise not less than $\mathbf{c} = (1, \ldots, n_1, 1, \ldots, n_2, \ldots, 1, \ldots, n_k)$. Clearly the existence of such ordering establishes the lemma. For $n_k = 1$ this is self evident. For $n_k > 1$ we apply 'greedy' approach – put all $a + 1$ (or a in case $n_k|n$) n_k's in the places of $n_k, n_{k-1}, \ldots, n_{k-a+1}$. The fact that $n_k \geqslant n_{k-1} \geqslant \ldots \geqslant n_1$ ensures that $k \geqslant a$ and all of $n_k - 1, n_{k-1} - 1, \ldots, n_{k-a+1} - 1, n_{k-a}, \ldots, n_1$ are less than n_k. Therefore we may apply inductive assumptions to these numbers thus finishing the proof of the lemma. □

Lemma 17 For every $\varepsilon > 0$ there exist $K \in \mathbb{N}$ such that if $n_1, \ldots, n_k \leqslant n/K$, where $n = \sum_{i=1}^k n_i$, then $\sqrt[n]{\prod_{i=1}^k n_i!/n!} < \varepsilon$.

Proof. Assume that $n_1 \leqslant \ldots \leqslant n_k \leqslant n/K$ and let $n = an_k + r$, where $0 \leqslant r < n_k$. By Lemma 48 we get that

$$\frac{\prod_{i=1}^k n_i!}{n!} \leqslant \frac{(n_k!)^a(n_k - r + 1)\ldots n_k}{n!} \leqslant \frac{1}{1^{n_k}} \cdot \frac{1}{2^{n_k}} \cdot \ldots \cdot \frac{1}{a^{n_k}} \cdot \frac{1}{(a+1)^r} \leqslant \frac{1}{(a!)^{n_k}}.$$

Therefore

$$\sqrt[n]{\frac{\prod_{i=1}^{k} n_i!}{n!}} \leqslant \frac{1}{\frac{n}{n\!\sqrt{a!}}} = \frac{1}{\frac{n}{n\!\sqrt{\lfloor \frac{n}{n_k} \rfloor !}}}.$$

For K large enough this might be arbitrarily small, so this finishes the proof. □

Lemma 49. *Let $(\mathcal{X}, d)$ be a finitely compact metric space, $D \subseteq \mathcal{X}$ a closed set and $(a_n)_{n=1}^{\infty}$ a bouded sequence in $\mathcal{X}$. If every converging subsequence of $(a_n)_{n=1}^{\infty}$ has a limit in D then* $\operatorname{dist}(a_n, D) \to 0$.

Proof. Suppose that $\liminf \operatorname{dist}(a_n, D) > 0$. Then there exist a subsequence $(a_{n_k})_{k=1}^{\infty}$ and $\varepsilon > 0$ such that $\operatorname{dist}(a_{n_k}, D) > \varepsilon > 0$. This contradicts the fact that $(a_{n_k})_{k=1}^{\infty}$ as a bounded sequence in $\mathcal{X}$ has a converging subsequence whose limit must belong to the closed set D. □

6.2 R Simulations

This subsection includes computer simulations concerning Example 45. The experiments were performed on a 64-bit Linux machine with `R` version 3.1.2 ([9]). To sample from the posterior distribution, Algorithm 2 from [8] was implemented. Sampling from multivariate normal and computing values of its distribution function were performed using `mvtnorm` package ([1]).

Firstly a random sample of 30 iid observations from Unif$[-1, 1]$ was sampled. Then 1000 MCMC samplings were performed, each starting from single cluster partition with 10 burn-in period. Finally the most frequent partition was chosen as MAP. The small number of observations is related to the fact that as their number grows, the MAP becomes increasingly difficult to distinguish among the sample. In case of 30 observations even if 100 samplings were performed, they produced different partitions. Therefore it is still very likely that the chosen partition is not the MAP but an approximation of it. It turns out, though, that even with such a small number of observations and an approximation of the MAP, our simulations still illustrate the theoretical results presented in Example 45. The simulations were performed for all possible combinations of $\alpha, \sigma^2, \tau^2 \in \{1, .1, .01\}$ (with the same vector of observation). The results are shown in Table 1.

The analysis of Table 1 yields to the following conclusions. Firstly, there is an apparent connection between the actual and predicted number of clusters in the MAP. Secondly, decreasing values of α and τ^2 leads to the smaller posterior probability of the chosen MAP consistent with the prediction basing on σ^2 and may even yield to the decrease of the number of clusters in the MAP. The impact of α is easily justified by the formula (1) as we have the punishment $\alpha^{\#\text{clusters}}$. The role of τ^2 is more difficult to explain as it occurs also in the formula for $R_m = \sigma^{-2} + \tau^{-2}/m$.

Of course with fixed vector of observation and decreasing α, the number of clusters in MAP tends to 1. However if we fix the parameters, e.g. set $\alpha = \tau^2 = \sigma^2 = 1$, and start to increase the number of observation, then the number of clusters in MAP slowly approaches our predictions. The results are presented in Table 2.

Finally, Figures 1 and 2 present illustrations of some of the partitions obtained. They correspond nicely to the results of convexity and equal segmentation, derived in Example 45.

α	τ^2	σ^2	$\mathbb{P}(MAP \mid \mathbf{x})$	occupation	est
1.00	1.00	0.01	0.013	8, 5, 5, 4, 3	5.77
1.00	0.10	0.01	0.007	8, 5, 5, 4, 3	5.77
1.00	0.01	0.01	0.098	12, 8, 3, 2	5.77
0.10	1.00	0.01	0.096	9, 6, 5, 5	5.77
0.10	0.10	0.01	0.053	9, 6, 5, 5	5.77
0.10	0.01	0.01	0.465	12, 8, 5	5.77
0.01	1.00	0.01	0.212	9, 6, 5, 5	5.77
0.01	0.10	0.01	0.185	12, 8, 5	5.77
0.01	0.01	0.01	0.789	12, 8, 5	5.77
1.00	1.00	0.10	0.017	15, 10	1.82
1.00	0.10	0.10	0.014	15, 10	1.82
1.00	0.01	0.10	0.009	16, 9	1.82
0.10	1.00	0.10	0.177	15, 10	1.82
0.10	0.10	0.10	0.121	15, 10	1.82
0.10	0.01	0.10	0.455	25	1.82
0.01	1.00	0.10	0.206	25	1.82
0.01	0.10	0.10	0.579	25	1.82
0.01	0.01	0.10	0.906	25	1.82
1.00	1.00	1.00	0.108	25	0.57
1.00	0.10	1.00	0.051	25	0.57
1.00	0.01	1.00	0.052	25	0.57
0.10	1.00	1.00	0.788	25	0.57
0.10	0.10	1.00	0.718	25	0.57
0.10	0.01	1.00	0.683	25	0.57
0.01	1.00	1.00	0.967	25	0.57
0.01	0.10	1.00	0.964	25	0.57
0.01	0.01	1.00	0.965	25	0.57

Table 1: Results of numerical experiments with all combination of $\alpha, \tau^2, \sigma^2 \in \{1, .1, .01\}$ and fixed $\mathbf{x}$. Column '$\mathbb{P}(MAP \mid \mathbf{x})$' shows the frequency of the most frequent partition in 1000 posterior sample. Column 'occupation' presents the occupation of tables created in MAP, sorted in decreasing order. Column 'est' is equal to $1/(\sqrt{3\sigma^2})$ and should approximate the number of clusters, cf. Example 45.

# added	$\mathbb{P}(MAP \mid \mathbf{x})$	occupation
25	0.409	20, 19, 11
50	0.177	30, 28, 17
75	0.089	36, 36, 28
100	0.019	43, 43, 39
125	0.025	53, 50, 47
175	0.015	73, 65, 62
225	0.004	94, 81, 75
275	0.002	107, 103, 90
400	0.001	120, 113, 98, 94

Table 2: Column ' # added' shows the number of observations added to initial 25 (i.e. a sequence of 1000 new observation was sampled and its prefixes of given length were added to the initial observations). The meaning of the other columns is the same as in Table 1.

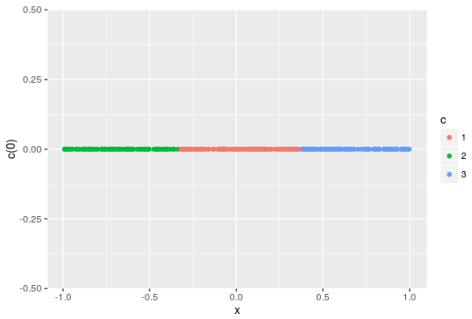

Figure 1: MAP of 300 observations on a segment

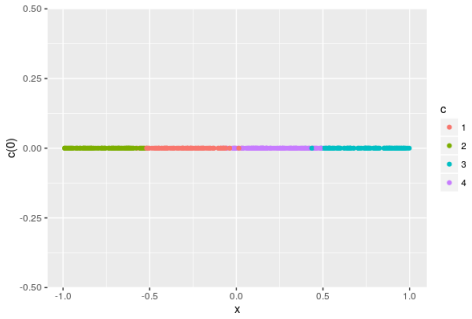

Figure 2: MAP of 425 observations on a segment

References

[1] Tetsuhisa Miwa Xuefei Mi Friedrich Leisch Fabian Scheipl Torsten Hothorni Alan Genz, Frank Bretz. mvtnorm: Multivariate Normal and t Distributions. R package version 1.0-5., 2016.

[2] Fuzhen Zhang (auth.). Matrix Theory: Basic Results and Techniques. Universitext. Springer-Verlag New York, 2 edition, 2011.

[3] J. L. Doob (auth.). Measure Theory. Graduate Texts in Mathematics 143. Springer-Verlag New York, 1 edition, 1994.

[4] William Feller. An Introduction to Probability Theory and Its Applications, Vol. 1, 3rd Edition. Wiley, 3 edition, 1968.

[5] D. Pollard J. Elker and W. Stute. Glivenko-Cantelli theorems for classes of convex sets. Advances in Applied Probability, 11, 12 1979.

[6] Jeffrey W. Miller and Matthew T. Harrison. Inconsistency of Pitman-Yor process mixtures for the number of components. Journal of Machine Learning Research, 15:3333–3370, 2014.

[7] Maria Moszyńska. Selected Topics in Convex Geometry. Birkhuser Boston, 1 edition, 2005.

[8] Radford M. Neal. Markov Chain sampling methods for Dirichlet Process mixture models. Journal of Computational and Graphical Statistics, 9, 06 2000.

[9] R Core Team. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria, 2014.

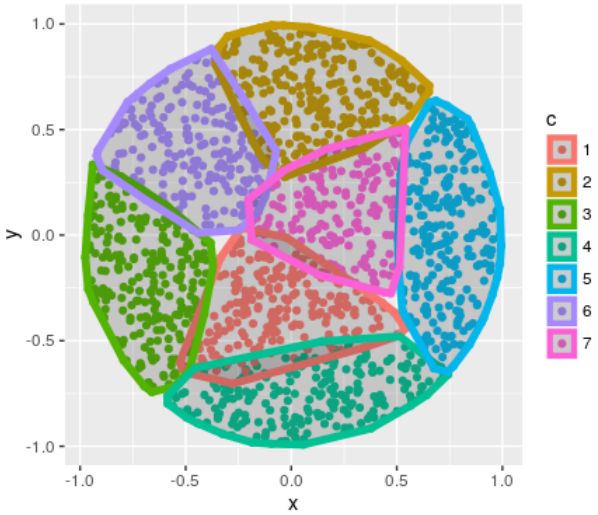

Figure 3: MAP of 1600 observations on a disc. The parameters of the model were $\alpha = .1$, $T = \mathrm{I}/10$ and $\Sigma = \mathrm{I}/100$.